



European Laboratory for Particle Physics
Laboratoire Européen pour la Physique des Particules
CH-1211 Genève 23 - Suisse

FEDKIT

FEDKIT user's manual and programmer's manual

Document Version: beta 0.9
Document Date: 12 November 2004
Document Status: Work version / subject to additions
Document Author: Eric Cano, Dominique Gigi

Abstract

This document introduces the use of the FED receiver and sender boards. Those boards and the tool kit are designed by the CMS data acquisition group for all the detector front end developers to provide read-out functionality from the Slink-64 connector on the FED.



1	Introduction	4
2	FEDKIT model	5
2.1	Hardware model	5
2.2	Software model	6
2.2.1	<i>Internal details of the driver</i>	6
3	FED receiver hardware documentation	8
3.1	Register list	8
3.2	Registers offsets	11
3.3	Software FIFO protocol	11
3.4	Word count FIFO records	12
3.5	Link test mode	12
3.6	Self send extension	12
4	FED software driver	13
4.1	Getting and installing the software	13
4.1.1	<i>Software requirements (kernel patch)</i>	13
4.1.2	<i>Retrieve the sources from CVS and compile</i>	13
4.1.3	<i>Install and link the software library, load the kernel modules</i>	14
4.1.4	<i>Example program</i>	15
4.2	API (programmer's manual)	16
4.2.1	<i>Initialisation functions</i>	16
4.2.2	<i>Fragment handling functions</i>	17
4.2.3	<i>Fragment analysis functions</i>	18
4.2.4	<i>Cleanup</i>	19
4.2.5	<i>Error handling</i>	20
4.2.6	<i>Suspending/restarting the receiver</i>	20
4.2.7	<i>Memory handling done by the user (optional behaviour)</i>	20
4.2.8	<i>Self send functionality</i>	21
4.2.9	<i>Link level dump functionality</i>	22
4.2.10	<i>FPGA version access</i>	22
4.3	Program behaviours in limit cases	23
5	Addition: FEDKIT sender (PCI to Slink64)	23
5.1	Hardware documentation	23
5.1.1	<i>Master mode</i>	23
5.1.2	<i>Slave mode</i>	24
5.1.3	<i>Generated event</i>	24
5.1.4	<i>FPGA revision number</i>	25
5.1.5	<i>Register summary</i>	25
5.2	Software documentation (API)	27
5.2.1	<i>Sending generated events</i>	27
5.2.2	<i>Using sender board as a slave</i>	28
5.2.3	<i>Using sender board as master</i>	28
5.2.4	<i>Getting the FPGA version</i>	29
6	Test programs	29
6.1	‘test_merge’: FEDkit’s main test program	29
6.2	‘fedkit_dump_receiver’: the link level debugging tool	30

- 7 Labview VIs 31**
- 8 Going further 32**
- References 32**



1 Introduction

In the CMS data acquisition, the event builder will retrieve event fragments from each sub-detector's front end driver (FED) through a normalised connector and protocol. This protocol was defined by the Readout Unit Working Group (RUWG) [1]. It is specified in the S-link 64 [2] protocol for the data transfer protocol and by the RUWG for the format of the event fragments.

The FEDkit development will be reused in the final DAQ input in the front end readout link board (FRL). Will be a plug and play replacement of the FEDkit in the final setup of CMS [3].

The CMS DAQ group is working on the design of boards and link that will take the data from the FED's Slink connector and send it through a cable to the receiving machine. The receiver board is based on the Generic III board (all purpose board based on an FPGA).

A few definitions to be clear with words: the CMS apparatus generates events, which are collected from the many FEDs (front end drivers) by the DAQ. The FEDkit implements this link from the FED to the DAQ. This document uses "fragment" to designate what comes out of the FED. On a more technical point of view, the FEDkit uses memory blocks to store the event fragments. So, from the FEDkit's point of view: a **event fragment**, or **fragment**, is made of one or more **blocks**.

The Generic III receiver board is a PCI 64 bit/66 MHz board that will sit on the receiver machine (PC). The sender board is also based on the GIII. Both boards communicate thanks to CMC (common mezzanine card) format daughter boards. The receiver board can receive data from one or two sender boards. The receiver board checks the validity of the CRC field for the fragment (CRC field is defined by the RUWG). Also, the receiver board can work in debug mode in which the words sent to the link are directly accessible to the software. This way a FED developer can debug the structure of fragments generated by the FED, if the data doesn't get through.

The blocks correspond to physical memory locations. The receiver board send the data to those locations by a DMA.

Finally, the FEDkit internally uses several FIFOs. Those FIFOs allow communication between the software and the hardware board. (this is detailed in Section 2, "FEDKIT model").

This document is still preliminary. Some informations are not known yet and are therefore printed in italic.

2 FEDKIT model

2.1 Hardware model

The FED kit is based on the link setup (two boards and cable). The receiver board sends the received data to the host PC's memory through the PCI bus.

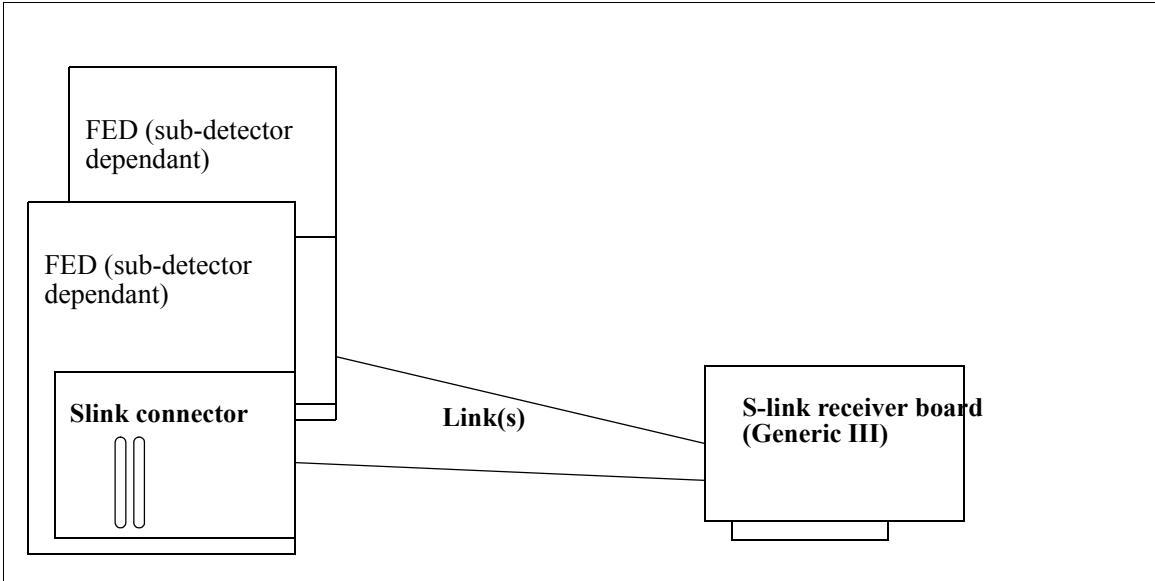


Figure 1 FED kit layout

The FED gets the L1 (hardware) triggers and sends automatically the event fragment to the S-link receiver. The model is a push model, but the S-link receiver throttles the data flow so that it won't overflow.

The receiver board works with DMA. It sends the received events through a DMA to memory blocks provided by the host. After DMAing an event fragment to host memory, the s-link receiver board simply sends the word count for the received event into a software FIFO (in memory, the word count FIFO). This scheme is developed in detail in Section 3.

When the receiver board starves on data blocks, it simply uses the back pressure feature of Slink. No overflow is therefore possible on the board (but this can create an overflow on the FED, which in turn is able to throttle the trigger rate). On the other hand, the board can ask for additional data blocks thanks to interrupts. This scheme was chosen to provide a arbitrary fragment size (it is limited by the memory blocks available and to 4GB (in practice less) because we work on 32 bits architecture).

Once the fragment has been fully transferred to the memory blocks, the receiver board notifies the software by posting on the word count FIFO the size of the transferred event. When this word count fifo gets full, the data flow is interrupted. The software knows which data blocks were passed to the receiver board and in which order, so knowing the size of the fragment, the software knows which blocks are used for each fragment.

The board can also be reset through a register. This function is used at initialisation time. The board has the ability to abort the transmission of a fragment in the middle of the transmission, but this feature is not currently used.

The receiver board also has the ability to dump all the words transmitted on the link, regardless of the fragment structure. Those functions, which also allow access to the control words can be used to debug another device generating fragments (a FED). The developer is then able to check what reaches the other side of the link. Only one link at a time can be debugged, as the merging uses the fragment structure.

An optional PCI sender/transmitter board is developed along with the receiver board for development. This board can send event fragments from PCI to the link. It can receive the fragments in slave mode or in DMA mode. Additionally, it can generate pattern events. As some FEDs won't be available at the beginning, this board could be of some use for the subsystems developers as well.

This document mainly focuses on the receiver board, but a chapter is devoted to the sender board. The API of the fedkit also covers the sender board. See Section 5, "Addition: FEDKIT sender (PCI to Slink64)" and Section 3.6, "Self send extension".

A second version of the receiver board receives the data from two links. The fragments from each FED are concatenated and make a single event fragment in the memory.

Later this merger will have some synchronisation features to handle the cases when the FEDs are out of sync, or when one of the 2 FEDs die (timeout).

Additionally, the data links (only one at a time) can be dumped to a PCI register. Therefore, the raw data on the link can be accessed. The data includes control bits and data payload.

2.2 Software model

The FED kit comes with complete driver that gives access to the functionality of the board. The user just has to rely on the provided driver to get the data coming from the FED. The basic software model is a set of C functions from which the user gets a new event fragment. Once the fragment data has been processed, the user frees the buffer with a function provided by the driver so that the driver can reuse the memory.

Additional functions are available to analyse the fragments (get data from header, get data from one FED in the case of a fragment coming from several FEDs).

2.2.1 Internal details of the driver

This part is provided for information only. The user of the fedkit could skip it.

The driver is a stand alone Linux driver. No additional driver is required, but the fedkit uses the bigphysarea memory to allocate its buffers. More details can be found in section

The driver is based on the toolkit (see [4]). This toolkit is useful to make device drivers that are multi-platform. Optionally, the fedkit can be compiled in Linux native version. This version contains a full featured device driver, which is much more robust to crashes of user processes.

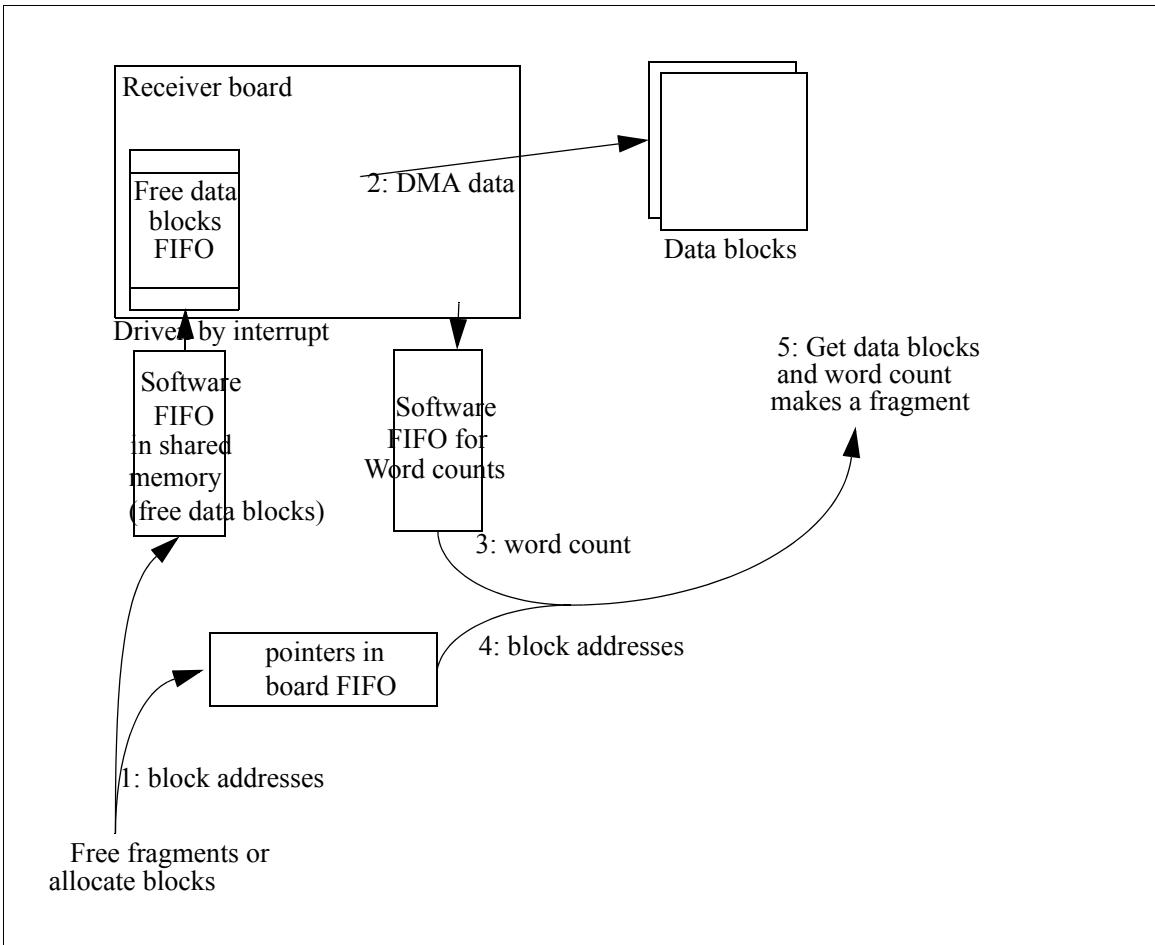


Figure 2 General layout of driver

The driver holds several lists inside the `fedkit_receiver` structure:

- free pointer FIFO: a list of the pointers to free blocks not already sent to the board in a software FIFO (as the board's free pointer FIFO is limited in size). This FIFO is store in a `bigphys` block, as it is filled by the user space functions of the fedkit, and emptied by the interrupt service routine (interrupt is triggered when the board's FIFO is half empty or empty). This FIFO uses $(n+1) * 4$ bytes in a single block of `bigphys` memory, where n is the number of blocks allocated by the fedkit at startup (this is 1000 by default). This FIFO is completely filled at initialisation, then the interrupt of the board is enabled, and the interrupt service routine start to feed the board's FIFO (as it is initially empty).
- pointers in board FIFO: a list of the pointers to free blocks sent to the board. This is just a chained structure in the driver. It is only used in user space. This chain is filled with the pointers to the same blocks as the free pointer FIFO, but is emptied by the process generating the fragment structure.

- word count FIFO: this FIFO is located in the host's memory. It is filled by the board itself, and emptied by the user space functions. When a word count is popped from this FIFO, the corresponding number of block addresses is popped from the pointers in board FIFO. This one is also stored in a bigphys block, and has a fixed size (it uses 2 blocks of 4kB on a x86 architecture).

All the bigphys blocks used here are actually Dbuffs (dma buffers of xdaq-shell, defined in the i2ocore package [4]) and accessible from user space, kernel space and PCI bus.

The interrupt is triggered when the hardware 'free block' FIFO is half empty or empty. The interrupt service routine then feed this FIFO with free block addresses from the pointers to free blocks FIFO.

Now when a fragment comes from the FED, but the board has no data blocks available, the software can send the "abort fragment" command. The board will then send the partial fragment to the word count FIFO, with a special flag. This feature is currently not used in the fedkit.

With this scheme, we can allocate an arbitrary number of blocks, limited only by the size of the bigphys memory. The data blocks themselves are also allocated from the bigphys memory itself.

3 FED receiver hardware documentation

The FED receiver board implements mainly a hardware FIFO to hold the free data blocks addresses and the support for a software FIFO in which word counts are sent.

Whenever it receives data from the Slink, the board DMAs data to the first available data block in its FIFO and updates an internal word count. If data spans over one block, the DMA engine goes on on the following block. When the free data block FIFO is empty, the process waits for a free data block to become available on the FIFO to go on. If the software driver is unable to provide free data blocks to the board, a special command (abort event) is used by the driver software.

As a separate process, the boards asks for more free data blocks when the FIFO is half empty and empty by generating interrupts.

The FED receiver board is a PCI board with vendor ID 0xECD6 and device ID 0xFD05.

3.1 Register list

This register list is summarized (and offsets are indicated) in Table 1.

- Block size register: write only (block size in bytes, has to be 64-bit aligned)
- DATA block FIFO (write only): PCI address of each block. The FIFO is 1024 addresses deep, and the driver implementor has to synchronise the sending of words with the FIFO half empty and FIFO empty interrupts
- Software FIFO registers (word counts - in bytes. Byte counts actually) see Section 3.3.
 - FIFO base address. Points to the PCI address of a 1024 word (4096 bytes zone)

- Read offset register (write only). Point to the next location to be read (not read yet).
- Write offset register address (where the board pushes the write offset pointer). Points to the first free location.
- Control and status register (CSR): read status write control.
 - The 8 LSBs indicate interrupt status. Interrupts are stateless, that is the only way to clear an interrupt is to remove the condition that triggers the interrupt. FIFO half empty is cleared by filling the FIFO. Interrupts can also be masked or enabled through the bits 8 to 15.
 - CSR[0]: reads 0 if FIFO is not empty. 1 if empty. An interrupt is triggered if CSR[8] is set to 1. CSR[0] is a read only bit. Interrupt is cleared by removing the triggering condition (or by disabling it).
 - CSR[1]: reads 0 if FIFO is more than half full. 1 if half empty. An interrupt is triggered if CSR[9] is set to 1. CSR[1] is a read only bit. Interrupt is cleared by removing the triggering condition (or by disabling it).
 - CSR[8]: enables the FIFO empty interrupt. This bit is readable and writable.
 - CSR[9]: enables the FIFO half empty interrupt. This bit is readable and writable.
 - The 16 most significant bits correspond to functions.
 - CSR[16] writing 1 triggers an abort event. Reads 1 if abort event is still in progress or 0 on the other case.
 - CSR[17] writing 1 start a “clear board” (software reset). Reads 1 if reset still in progress, 0 otherwise.
 - CSR[18] simple read/write bit. writing 1 enables the links (this includes a synchronisation operation over the links). Writing 0 disables all the link. Reading the bit gives the state of the link (enabled/disabled)
 - CSR[20] read only bit. When in link test mode, this bit indicates that at least an error was detected.
 - CSR[21] simple read/write bit. This bit selects whether the LINK_ERROR_BITS register displays the error results for the bits 0..31 (CSR[21]=0) or for 32..63 (CSR[22]=1).
 - CSR[22] simple read/write bit. Setting this bit to 1 will start the link test mode. The tests runs continuously as long as the bit is set. The result of the test is accessible through CSR[20], CSR[21]. While the test runs, the general operations of the FEDkit is only suspended. When the bit is set back to 0, the FEDkit operations continue from where they were suspended.
 - CSR[24] simple read/write bit. Writing 1 enables the first receive link. Writing 0 disable the link. This bit is to be used for fine grained channel selection. Defaults to 0.
 - CSR[25] simple read/write bit. Writing 1 enables the second receive link. Writing 0 disable the link. This bit is to be used for fine grained channel selection. Defaults to 0.



- CSR[28] simple read/write bit. Writing 1 enables the dumping of the link. In this mode, only one link can be selected in CSR[24,23] during debug operations. Writing back 0 re-enables the normal operations.

This mapping is summarised in Figure 3.

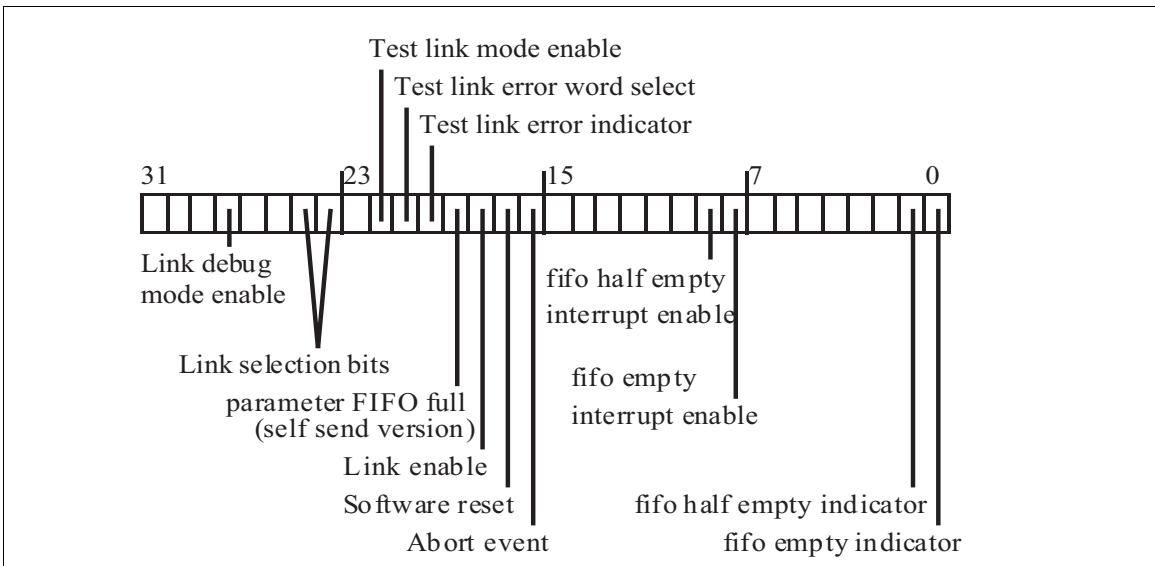


Figure 3 Bit mapping in the control word

- Link level dump registers. Before using those, user has to set the debug bit (CSR[28]). Also only a single link can be used (CSR[24,25]). The user has to read the 4 registers DUMPDATA0, DUMPDATA1, DUMPCTRL0, DUMPCTRL1 before those registers move to the next word on the link.
 - By monitoring FIFOSTAT, the user can tell whether there is something to read in the PCI FIFO (FIFOSTAT[2]). FIFOSTAT[0,1] respectively indicate that the FIFOs for link 0 and 1 are non empty. Those bit are indications only to the software, as there is no direct access to those values from PCI.
 - DUMPDATA0, DUMPDATA1 respectively contain least and most significant 32-bits words from the current 64-bits word from the link.
 - DUMPCTRL0 only contain a single bit (DUMPCTRL0[0], which reflects the state of the control bit on the line for this word.
 - DUMPCTRL1 currently contain no data.
- Link test mode result register. The link test mode is controlled from CSR[20..22]. The LINK_ERROR_BITS register indicates which bits were affected by at least one error.
- On offset 48h of the configuration space, a read-only register indicates the 32-bits firm ware revision number.

3.2 Registers offsets

Table 1 Registers description

Mnemonics	Register description	Offset from base address 0
BKSZ	Block size (in bytes, must be 64 bits aligned)	80h
WCFADDR	Word count FIFO address	84h
WCFR	Word count FIFO read offset (in words) see Section 3.3	08h
WCFWADDR	Word count FIFO: address of the write register address in memory	88h
FRBK FIFO	Free block FIFO	04h
CSR	Control and status register	00h
FIFOSTAT	FIFOs status	110h
DUMPDATA0	Bits 31..0 of the data payload	100h
DUMPDATA1	Bits 63..32 of the data payload	104h
DUMPCTRL0	Control bits of the link	108h
DUMPCTRL1	Reserved (used to control fifo read)	10Ch
LINK_ERROR_BITS	Bits 31..0 of the error pattern in link test mode if CSR[21]=0. Bits 63..32 otherwise.	180h

3.3 Software FIFO protocol

The offset variables (read and write point to the next word to be written (or read)). Every time there is a write or read, the corresponding variable is incremented. Those variable are eleven bit variables. The 10 least significantly bits (bits [9..0])contain the number of the block to be next read or written. The additional MSB (bit 10) is added to the number. When a write is done, this variable is increased by one (and wrap to 0 after 2048 increases).

This trick allows us to know if the FIFO is completely empty or completely full when write[9..0] == read[9..0]

If write [10] == read[10] (and write[9..0] == read[9..0]), the FIFO is empty (we can write, but not read) the other way round, if write[10]!=read[10] (and write[9..0] == read[9..0]), the FIFO is full (we can't write, but we can read). The board and the driver keep a copy of each value, so FIFO can work. The software has to poll the FIFO to know if it's empty or not.

3.4 Word count FIFO records

The 32 bits words in the word count FIFO are organised as follows (see Figure 4):

- Bit 31: Indicates that the receiver detected a CRC error (CRC is defined as in [1]).
- Bit 30: Indicates that receiver had to truncate the fragment because its size would not fit in the 24 bits of the word count (indicated the truncation of a fragment bigger than 16MBytes)
- Bits 23-0: word count (in bytes)

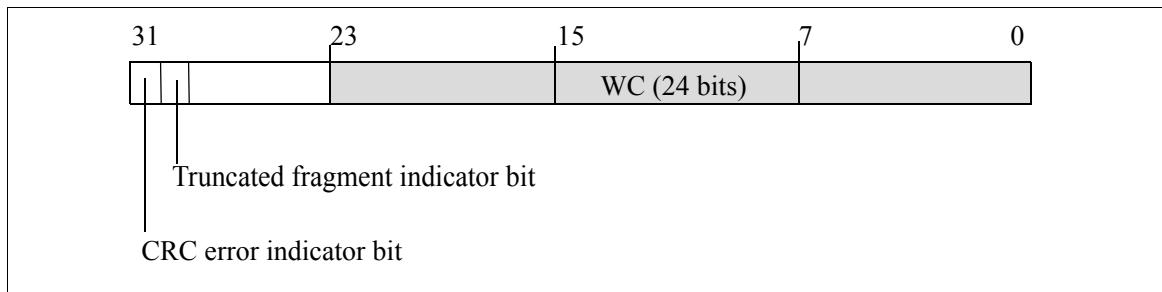


Figure 4 Bits assignment in word count FIFO records

3.5 Link test mode

The FEDkit receiver hardware can run tests of the link cable independently of the software. To run in this mode, the

3.6 Self send extension

A special version of the board was developed for software design. Instead of requiring a sender and receiver setup, this special receiver board has a generator interface on another PCI base address. This allows simpler hardware setups for developing receiver software setups. A special function supporting this feature is included in the fedkit receiver API. (see Section 4.2.8, "Self send functionality")

In a possible extension the board could generate an interrupt when the word count FIFO is empty. The driver would then rely on interrupt and not on polling to know when an event is available.

Another additional extension could be to have the receiver board writing the word count plus a pointer address (last or first pointer, this doesn't matter) for each event in order for the software to detect if software FIFO and hardware FIFO are still in sync (they should be all the time, so this is just for additional robustness). This would be hidden from the user by the fedkit software).

4 FED software driver

4.1 Getting and installing the software

The latest release of the software, the latest documentation, etc... can be found on the fedkit web page:
<http://cern.ch/cano/fedkit>.

4.1.1 Software requirements (kernel patch)

The FEDKIT needs the bigphysarea patch to be installed. The version of this patch matching the Linux kernel you are using can be found on the page:

<http://www.polyware.nl/~middlelink/En/hol-v41.html>

This kernel sources can be retrieved from AFS:

`/afs/cern.ch/project/linux/dev/kernel/2.4-current/build-current`

Then `/etc/lilo.conf` just has to be updated to boot from this image. Also add as a kernel parameter the size of the memory to be reserved for bigphys:

`append="bigphysarea=16384"`

For grub, simply add `bigphysarea=16384` at the end of the kernel line.

4.1.2 Retrieve the sources from CVS and compile

The sources are stored in a CVS repository. They can be retrieved from the CMS CVS server the following way:

```
# cvs -d :pserver:anonymous@cmscvs.cern.ch:/cvs_server/repositories/TriDAS login
(password for anonymous checkout is "98passwd")

# cvs -d :pserver:anonymous@cmscvs.cern.ch:/cvs_server/repositories/TriDAS co \
    -P -r fedkit_v_1_30 TriDAS/daq/itools \
    TriDAS/daq/extern/i2o
```

The tag (here `fedkit_v_1_20`) might change over time. It is not recommended to get the head as you could get development versions. The various versions of the software and the various tags can be found on the web page of the fedkit [6] as well as all the revisions of this document.

Now the default version of the fedkit is the Linux native one as the only use of the fedkit is done on Linux systems.

The fedkit library is compiled by doing:

```
#make fedkit
```



in the TriDAS/daq/itools directory. This will create a package subdirectory containing all the libraries and the include files.

To install the drivers, the user has to do:

```
#make fedkit-install
```

You need either to be a sudo-enabled user or to do this as root. Once this has been done, the user doesn't have to load the driver as the driver will load automatically when opening the fedkit file.

The old i2ocore version of the driver is still available. Its use is not advised.

To compile the i2ocore version, just do:

```
#make fedkit-i2ocore
```

The API of the fedkit is the same in both cases (and many functions are strictly the same in both cases).

4.1.3 Install and link the software library, load the kernel modules

4.1.3.1 Using the i2ocore (obsolete) version

To compile a program using fedkit, the header file “fedkit.h” has to be included, and the final executable has to be linked with `libfedkit.a` and `libxdaq-shell.a`.

The required includes are in the directories

```
TriDAS/daq/itools/packages/fedkit/include  
TriDAS/daq/itools/packages/i2o/include  
TriDAS/daq/itools/packages/xdaq-shell/include
```

and libraries are:

```
TriDAS/itools/packages/fedkit/libfedkit.a  
TriDAS/itools/packages/xdaq-shell/libxdaq-shell.a
```

To run the software, 2 special files have to be created:

```
crw-rw-rw- 1 root      root      124,      0 Dec 14 15:03 /dev/xdaq-shell  
crw-rw-rw- 1 root      root      124,      1 Nov 22 11:47 /dev/xdaq-shell-pmap
```

The following commands create those files (as root, or using sudo):

```
/bin/mknod /dev/xdaq-shell c 124 0  
/bin/chmod a+rw /dev/xdaq-shell  
/bin/mknod /dev/xdaq-shell-pmap c 124 1  
/bin/chmod a+rw /dev/xdaq-shell-pmap
```

or the Makefile can be used (as root, or using sudo):

```
make device-files
```

Then before running the software, the kernel modules have to be loaded (as root, or using sudo):

```
/sbin/insmod packages/i2o/i2ocore.o  
/sbin/insmod packages/xdaq-shell/xdaq-shell.o  
/sbin/insmod packages/fedkit/fedkit_kernel.o
```

They can also be loaded through the main Makefile (as root, or using sudo):

```
make load-fedkit
```

4.1.3.2 Using the Linux native version

As explained in the previous part, the native version doesn't require i2ocore and xdaq-shell library. The user doesn't need to link to libxdaq-shell.a.

Note that the fedkit-native test program (test_merge) still makes use of the xdaq-shell to allocate some DMA buffer. Therefore, when making fedkit, the xdaq-shell is still compiled. The user also has to install xdaq-shell in order to use the noalloc mode of the receiver or the master DMA mode of the sender in the test_merge program. All other modes don't use the DMA buffers.

The xdaq-shell driver is also installed by the make fedkit-install described above.

4.1.4 Example program

An example program is provided with the library. The file is packages/fedkit/fedkit-example.c. It can be compiled by executing the following line from the TriDAS/itools directory:

```
gcc packages/fedkit/fedkit-example.c -o packages/fedkit/fedkit-example \
-I packages/fedkit/include -I packages/xdaq-shell/include \
-I packages/i2o/include -D BIG_ENDIAN_ \
packages/fedkit/libfedkit.a
```

4.2 API (programmer's manual)

4.2.1 Initialisation functions

The first task of the user is to connect with a FED receiver board (fedkit_open). The software can handle several boards in the same system. The software keeps track of each board through a descriptor (struct fedkit_receiver) provided by the fedkit_open function.

```
struct fedkit_receiver * fedkit_open (int index, int *status);
```

This function returns NULL in case of failure (board not found). The reason for the failure is then coded into the status integer. If the user provides a NULL pointer in this place (meaning "don't care") no error code will be available. This function can be used to scan for all the receivers present in the system, testing from index 0 onward. When the user reaches an index number unavailable, he knows he found all the boards.

The user can then set the required parameters for the board. The parameters are:



1. Block size (the default size is *TBD*). For better efficiency, the block size should be a multiple of the architecture's page size (4kbyte on Intel). This size has to be 64 bits aligned.
2. The number of blocks the FEDKIT should allocate. In this version, all the blocks have to be allocated at the beginning. This step is required. The block number has to be at least 2k (2048) as the interrupt algorithm relies on having enough blocks in circulation.
3. Header reserved size in bytes (the default is 0). The header size has to be aligned on a 64 bit boundary because of the DMA (64 bit DMA).
4. Link selection. this function is available only on the merger versions of the board. When using a merger board, the user can select which links are active. The pattern set by the function is just a binary pattern: bit 0 set to 1 to enable link 0, bit 1 for link 1, etc... (in the current design, we have 2 links maximum). By default, the software sets the pattern to 1 (only link 0 enabled).
5. *A footer size could be added in the future with the same alignment requirements.*

The 'set' function return positive or null parameter value, or a negative number on error (error codes are *TBD*). The 'get parameter' functions return positive or zero parameter value on success and negative error code on failure.

Functions prototypes are:

```
int fedkit_set_block_size (struct fedkit_receiver *receiver, int block_size);
int fedkit_get_block_size (struct fedkit_receiver *receiver);

int fedkit_set_block_number (struct fedkit_receiver *receiver,
                           int block_number);
int fedkit_get_block_number (struct fedkit_receiver *receiver);

int fedkit_set_header_size (struct fedkit_receiver *receiver, int header_size);
int fedkit_get_header_size (struct fedkit_receiver *receiver);

int fedkit_set_receive_pattern (struct fedkit_receiver *receiver, int pattern);
int fedkit_get_receive_pattern (struct fedkit_receiver *receiver);
```

Once those parameters have been set, the driver has to be started. When the board has been started, no parameter can be changed. (Some blocks are already in the FIFO of the board). The start function returns 0 on success or negative error code on failure.

```
int fedkit_start (struct fedkit_receiver *receiver);
```

Once the receiver board has been started, it is ready for reception (the driver automatically allocates memory blocks and feeds them to the board).

4.2.2 Fragment handling functions

To get the received event fragments, the user just has to get the received events, and to release them after processing. The event metadata (pointers to blocks, *possibly event ID and various parameters*) is stored in the fedkit_event structure. The first step is to get a fragment:

```
struct fedkit_fragment * fedkit_frag_get (struct fedkit_receiver * receiver,
                                         int *status);
```

The ‘get fragment’ function returns NULL if no more event has been received, or in case of error. The error reason is encoded in the status integer (if the `status` pointer was not NULL, which is a valid entry meaning “don’t care”). The reason can be “no event available”.

Several informations about the fragment can be extracted from the structure. The user can query the number of blocks, the block size (which can be different from receiver to receiver), the header size, and the size of the fragment. As in previous cases, getting a negative value means an error condition.

```
int fedkit_frag_size (struct fedkit_fragment *fragment);  
int fedkit_frag_header_size (struct fedkit_fragment *fragment);  
int fedkit_frag_block_size (struct fedkit_fragment *fragment);  
int fedkit_frag_block_number (struct fedkit_fragment *fragment);
```

Everything is quite self explanatory, beside the fragment size, which is the size of the payload only (the sizes of the headers are not included).

Once the fragment object is acquired, the user can access directly the blocks from the payload. The pointer is a `void *`, pointing to the start of each block (hence start of the header) and the user has to cast it to the desired structure or data type.

The blocks can be accessed through the function `fedkit_frag_block`:

```
void * fedkit_frag_block (struct fedkit_fragment *, int number);
```

Finally, once the fragment has been processed, the memory can be returned to the driver:

```
void fedkit_frag_release (struct fedkit_fragment *);
```

The memory for this fragment will be recycled into the board.

The hardware automatically checks a fragment’s CRC on the fly. It also indicates if it had to truncate a fragment which is larger than the maximum size it can handle (currently 16MBytes). Those function return non-zero in case of CRC error, or fragment truncation, respectively:

```
int fedkit_frag_get_CRC_error (struct fedkit_fragment * fragment);  
int fedkit_frag_get_truncated (struct fedkit_fragment * framgent);
```

4.2.3 Fragment analysis functions

In order to be able to analyse easily fragments coming from the fedkit, and especially from the merge version (2 fragments from 2 FEDs are concatenated in a single fragment by the board). Those functions analyse the structure of the fragment. Specifically, the RUWG defined a field on the trailer word of the FED payload that contains the size of the fragment. This field is used to recover all the FED fragments in the fedkit fragment (in the case of the merge). Those functions make some sanity checks on the fragment, and fail if the fragment structure is not right. The index used on all those functions is 0 for the first FED’s part and 1 for the second FED’s part.

The functions are the following:

```
int fedkit_frag_get_FED_number (struct fedkit_fragment * fragment);
```

this function returns the number of FED fragments in the fragment (this is useful in the case of the merge, or negative (-error) if there was a problem.

```
void * fedkit_frag_get_FED_trailer (struct fedkit_fragment * fragment,  
                                     int FED_index);
```

This function returns a pointer to the trailer word of the given FED fragment, or NULL in case of failure.

```
void * fedkit_frag_get_FED_word (struct fedkit_fragment * fragment,  
                                   int FED_index, int index);
```

This function returns a pointer to a given word in the FED fragment, or NULL in case of failure.

```
U32 fedkit_frag_get_FED_word_pci (struct fedkit_fragment * fragment,  
                                   int FED_index, int index);
```

This function returns the PCI address of a given word in the FED fragment, or 0 in case of failure. This function was mainly useful in the process of debugging the FEDKIT.

```
int fedkit_frag_get_FED_trigger (struct fedkit_fragment * fragment,  
                                 int FED_index);
```

This function extracts the positive trigger number (LV1_ID) from the FED fragment. Returns negative error in case of failure.

```
int fedkit_frag_get_FED_wc (struct fedkit_fragment * fragment, int FED_index);
```

This function extracts the positive word count from the FED fragment. Returns negative error in case of failure.

```
int fedkit_frag_get_FED_block_number (struct fedkit_fragment * fragment,  
                                       int FED_index);
```

This function returns the number of blocks the payload (Slink64 header and trailer excluded) of the given FED's part of the fragment spans onto. Returns negative (-1) in case of error, can return 0 if the payload is empty.

```
void * fedkit_frag_get_FED_block_pointer (struct fedkit_fragment * fragment,  
                                         int FED_index, int index);
```

This functions returns a pointer to the nth block of the given FED's part in the fragment (if index is not right, returns NULL);

```
int fedkit_frag_get_FED_block_size (struct fedkit_fragment * fragment,  
                                    int FED_index, int index);
```

This function returns the number of 64 bits words of payload contained in the block number 'index' of the FED number 'FED_index' in the given fragment.

Many functions allow access to the various fields from the headers and trailers of a given fragment. They all return a negative error value if there is a problem accessing the field, or the value of the field in the other case (every field is smaller than 32 bits). The functions are:

```
int fedkit_frag_get_FED_eventID (struct fedkit_fragment * fragment,  
                                 int FED_index, int index);
```

```
int fedkit_frag_get_FED_trigger (struct fedkit_fragment * fragment,  
                                 int FED_index);
```

```
int fedkit_frag_get_FED_BX_id (struct fedkit_fragment * fragment,
```

```
        int FED_index);  
  
int fedkit_frag_get_FED_LV1_id (struct fedkit_fragment * fragment,  
                                int FED_index);  
  
int fedkit_frag_get_FED_Source_id (struct fedkit_fragment * fragment,  
                                    int FED_index);  
  
int fedkit_frag_get_FED_Evt_lgth (struct fedkit_fragment * fragment,  
                                    int FED_index);  
  
int fedkit_frag_get_FED_Evt_stat (struct fedkit_fragment * fragment,  
                                    int FED_index);  
  
int fedkit_frag_get_FED_CRC (struct fedkit_fragment * fragment,  
                             int FED_index);  
int fedkit_frag_get_FED_FOV (struct fedkit_fragment * fragment,  
                             int FED_index);  
int fedkit_frag_get_FED_Evt_type (struct fedkit_fragment * fragment,  
                                    int FED_index);
```

4.2.4 Cleanup

As a last action, the board has to be released by the drivers, which will release all the resources used:

```
void fedkit_close (struct fedkit_receiver *);
```

This function can be called at any time.

4.2.5 Error handling

To ease debugging and error reporting, the function:

```
const char * fedkit_get_error_string (int error_number);
```

allows the user to convert error codes into more meaningfull messages for logging, error display, etc...

4.2.6 Suspending/restarting the receiver

The following functions are mainly useful to suspend the receiver so that a list of the blocks internally used can be retrieved (as described in Section 4.2.7, "Memory handling done by the user (optional behaviour)").

The suspend function is:

```
void fedkit_suspend (struct fedkit_receiver * receiver);
```

and the resume function:

```
void fedkit_resume (struct fedkit_receiver * receiver);
```



4.2.7 Memory handling done by the user (optional behaviour)

In addition to the schemes proposed in the previous paragraphs, the user can also provide the fedkit with his own memory blocks. Those blocks have to have an address in kernel space, user space and memory space. The blocks provided by the Dbuffs (DMA buffer from i2ocore [4]) are an option for this task. The first thing to do is to use the alternative function:

```
int fedkit_start_noalloc (struct fedkit_receiver *receiver);
```

This function will do exactly the same work as the normal version, but will not allocate the data blocks. Also, the software FIFO that holds the blocks has a fixed size after `fedkit_start` functions have been called. Therefore, the user should make sure to have set properly the block number parameter. The next functions to be used are:

```
int fedkit_provide_block (struct fedkit_receiver * receiver,
                           void *user_address, void *kernel_address,
                           void* bus_address, void *user_handle);
```

The `user_handle` parameter is a free pointer that will be kept with the data block. The user will typically store a pointer to his own structures to keep track of the block. The user can recover this pointer for each of the blocks allocated in a fragment by using:

```
void * fedkit_frag_user_handle (struct fedkit_fragment * fragment, int index);
```

When the noalloc mode has been chosen, additional functions can be used.

The function `fedkit_frag_release` still recycles the data blocks, but the new function

```
void fedkit_frag_release_norecycle (struct fedkit_fragment * fragment);
```

will not recycle the blocks (so it's the user's duty to call `fedkit_provide_block` again for those blocks).

When the user call `fedkit_close`, no block is freed, so the user has to keep track of his blocks for final deallocation.

Another scheme can be used by the user to recover the handles of all the blocks that have been provided to the board. Before doing this, the user has to suspend the board. Once the board has been suspended, the user can restart it (using `fedkit_suspend/fedkit_resume`).

Once the board has been suspended, the `fedkit_frag_get` function return NULL event if an event is pending in the memory. (So that the blocks stay held by the receiver and are not returned to the user).

The first function to be called is:

```
int fedkit_susp_get_handle_number (struct fedkit_receiver * receiver);
```

This function indicates the number of handles (actually of blocks) stored in the hardware or software FIFOs of the board. This function fails (by returning negative number) if the board is not suspended.

Then the user can recover the handles for each block by calling as many time as necessary the function:

```
void * fedkit_susp_get_handle (struct fedkit_receiver * receiver,
                               int index);
```

with index going from 0 to `fedkit_suspended_get_handle_number (receiver) -1` included. This function returns NULL if an invalid index is provided or if the receiver is not suspended.

After using those functions, `fedkit_close` can be called at any time. It should be noted that the blocks for which a handle has been retrieved can still be used by the FEDkit, so the user must call `fedkit_close` before reusing the blocks in another application. The best way to operate is to recover all the handles in a temporary array, call `fedkit_close`, and then only, reuse or free the blocks.

4.2.8 Self send functionality

A special version of the receiver board (special FPGA programming) generate data locally (see Section 3.6, "Self send extension"). This extension is automatically supported in the fedkit API. There just is an additional function, that fails if the hardware function is not present. This function also sends back an error when the parameter FIFO is full. (No space left to send parameters for events to be generated).

The function is:

```
int fedkit_autosend_generated (struct fedkit_receiver * receiver,  
                               U16 word_count, U16 seed, U32 event_trigger_number);
```

An additional function is provided to detect if a given receiver is a normal receiver or simply a self-send version.

```
int fedkit_is_autosender (struct fedkit_receiver * receiver);
```

returns 0 if the receiver is a normal one and non-zero if it is a self sender.

4.2.9 Link level dump functionality

The latest revisions of the FEDkit now include a functionality allowing the user to dump the raw data coming to the FEDkit receiver, regardless of the fragment structure. Using those functions allows the user to access this low level data.

The fedkit package also provides the user with a `fedkit_dump_receiver` program that simply dumps the data to the standard output until killed. Try `fedkit_dump_receiver --help` to check the available options.

To use the dump mode in a program, the user first has to call:

```
int fedkit_enable_link_dump (struct fedkit_receiver * receiver);
```

This function will return `FK_OK` in case of success, otherwise in case of error. This operation suspends the normal operation of the FEDkit.

Then the words are dumped one by one using

```
int fedkit_link_dump (struct fedkit_receiver * receiver,  
                      U32 * lsw, U32 * msw, int* control);
```

This function return `FK_OK` in case of success, `FK_error` if something goes wrong (link dump mode not enabled), or `FK_empty` if there is nothing to read in the FIFOs (nothing received from the link). In case of success, the least significant 32 bits word of the 64 bits word read from the link will be put in

*lsw. The most significant in *msw, and *control will be non-zero if and only if the control bit was set for this word.

Finally, to return to the normal mode of operations, the user has to call:

```
int fedkit_enable_link_dump (struct fedkit_receiver * receiver);
```

4.2.10 FPGA version access

In the latest versions of the FPGA, a unique ID is assigned to each revision of the firm ware. This ID is accessed through the following function:

```
U32 fedkit_get_FPGA_version (struct fedkit_receiver * receiver);
```

4.3 Program behaviours in limit cases

The driver gracefully frees all the allocated memory (including bigphys memory) event in case of a segmentation fault. In the native version of the FEDkit driver, a new program can use the hardware immediately, event if the previous program accessing the FEDkit crashes.

5 Addition: FEDKIT sender (PCI to Slink64)

During the prototyping activities, the FEDs won't be available, so in order to feed the Slink receiver, a sender board was developed. The FEDKIT sender board is a PCI 64/66 board (small PCI form factor) able to send data over the Slink64. The data can be generated by the board, send to the board (the FEDKIT sender is then a PCI slave), or the FEDKIT can pull the data on the PCI bus (DMA).

5.1 Hardware documentation

The board is a PCI board based on the generic III FPGA board. Vendor ID is 0xECD6 and device ID is 0xFE01.

The boards works in three modes:

- Pseudo random generator
- Slave
- Master

5.1.1 Master mode

In master mode, the user has to send the PCI address of the data location on base address 0, offset 0, then write the event size to base address 0, offset 4. The DMA will start immediately after the size has been written. The user has to send parameters for one event at a time, i.e. address/size then next address/next size and so on.

To check if the DMA is finished, the user can resort either to polling or can wait for an interrupt. The polling and interrupt both are controlled from the master control word (MCTRL) at base address 0, offset 0x8. This word contains the “DMA done” interrupt status bit (bit 3) the interrupt enable bit (bit 11), the “DMA done” interrupt acknowledge (bit 17) and the “DMA in progress bit” (bit 18).

- Bit 3 reads 1 when an interrupt is triggered, 0 otherwise. Interrupt is cleared by writing 1 to bit 18.
- Bit 11 is a read/write bit and writing 1 enables the end of DMA interrupt.
- Bit 18 is a status bit indicating if a DMA is in progress (read 1, 0 otherwise). This bit is to be used for polling.

The event must contain the control words as first and last word, i.e. the first word must be 0x5xxx.xxxx.xxxx.xxxx and the last word must be 0xAxxx.xxxx.xxxx.xxxx.

In order to use the master mode, the board has to be put in master/slave mode, by setting the bit 16 of MCTRL to 0. (Setting this bit to one puts the board in generator mode).

5.1.2 Slave mode

In slave mode, the user has to signal start and end of event fragment to the board. The start of event fragment is signalled by writing a 64 bit word of the format 0x5xxx.xxxx.xxxx.xxxx in the slave control register (base address 1, offset 0). Then the user writes as many data words as needed to base address 1, offset 8 and following. The end of the event is signalled by writing a word of format 0xAxxx.xxxx.xxxx.xxxx to the slave control register.

The maximum size of a burst is 64k bytes. To send a bigger event, the user has to restart a second DMA with the remaining data to offset 0x8 and following.

To use the slave mode, the user has to put the board in master/slave mode (MCTRL bit 16 to 0).

5.1.3 Generated event

Sender board can generate events following the RUWG specifications [1].

Generated mode is controlled by a control and status register (GCTRL) located at base address 2, offset 0.

In order to use the generated event mode, the board has to be put in generated mode (set GCTRL bit 16 to 1).

In generated mode, the user writes three times to base address 2, offset 0x4. The first word is the Event trigger number (24 bits), second word is the 64bits-word count (25 bits), third word contains source number (11 bits in bits 27..16), or-ed with the 8 bits seed for random data (8 least significat bits). The source number is only supported since version 0x3F200009 of the firmware. Those words have to be written in sequence. As soon as the last word (source number/seed) is written, the generator start to send the generated event over the link. The board can contain the parameters for several events in a FIFO. The user can control how many events he can send to the FIFO by polling the indicators FIFO empty (GCTRL, bit 0), FIFO half full (GCTRL, bit 1) and FIFO full (GCTRL, bit 2), or by relying on the interrupts generated by those status words (those interrupts are enabled by the respective bits 8,9 and 10).

5.1.4 FPGA revision number

On offset 48h of the configuration space, a read-only register indicates the 32-bits firm ware revision number.

The source number support in generated mode was supported started with version 0x3F200009 of the FPGA. It is ignored in previous version.

5.1.5 Register summary

The following figure summarises the bit map for the control register (MCTRL/GCTRL)

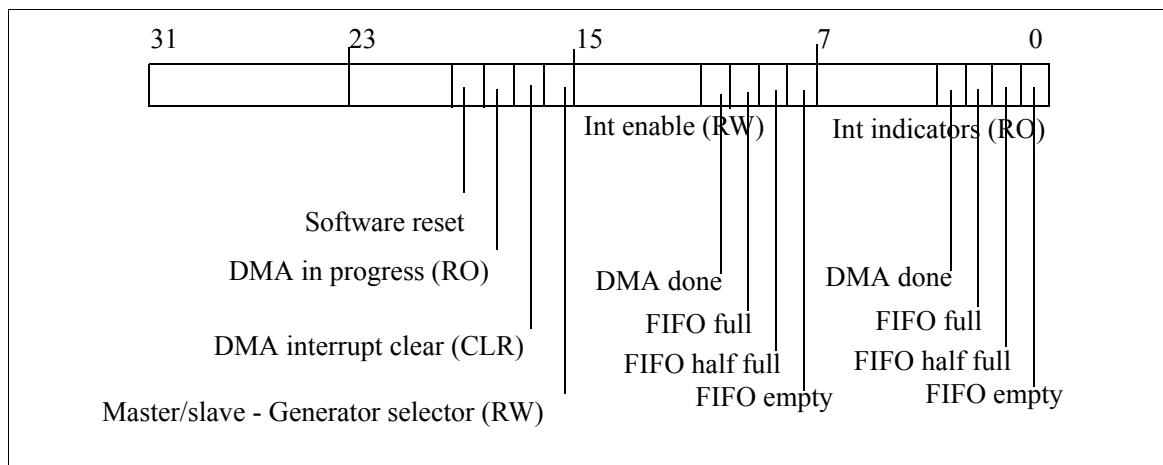


Figure 5 Bit map summary for MCTRL/GCTRL

The following table summarises the various offsets of the sender board.

Table 2 Register map for the FEDKIT sender

Base address	Offset	Register
Configuration space	0x48	FPGA revision number (read-only)
0	0x0000	MEVTADDR: Master mode: write address of event in memory. This write should be followed by a write to MEVTSZ. The first word of the event fragment should be 0x5xxx.xxxx.xxxx.xxxx and the last word should be 0xAxxx.xxxx.xxxx.xxx
0	0x0004	MEVTSZ: Master mode: write event size (in 64 bit words). Writing event size starts the DMA.
0	0x0008	MCTRL: Master mode and generated control word. <ul style="list-style-type: none"> Bit 0 indicates that the parameter FIFO of the generator is empty. This bit is read only. Bit 1 indicates that the parameter FIFO of the generator is half full. This bit is read only. Bit 2 indicates that the parameter FIFO of the generator is almost full. The FIFO cannot hold the parameters for an additional event fragment. This bit is read only. Bit 3 indicates that a DMA (master mode) has reached completion. This bit is read only Bit 8 is the FIFO empty interrupt enable. An interrupt is generated when this bit is one and the FIFO is empty. This bit is readable and writable. Bit 9 is the FIFO half full interrupt enable. An interrupt is generated when this bit is one and the FIFO is half full. This bit is readable and writable. Bit 10 is the FIFO full interrupt enable. An interrupt is generated when this bit is one and the FIFO is full. This bit is readable and writable. Bit 11 is the DMA done interrupt enable. An interrupt is generated when a master mode DMA reaches completion. This bit is readable and writable. Bit 16 is the selector for master/slave mode or generator mode. When it is set to 0, the board works in master/slave mode. When it is set to 1, the board works in generator mode. This bit is readable and writable. Bit 17 DMA done interrupt clear. This bit always reads 0. Writing one to it clear the DMA done interrupt (and bit 3), if there is a DMA done interrupt pending. Bit 18 is the DMA in progress bit (master mode). This read only bit reads 1 when a master mode DMA is in progress, 0 otherwise. Bit 19 is the software reset. Writing 1 to it makes the board go back to its default state. This bit always reads back 0.



Table 2 Register map for the FEDKIT sender

Base address	Offset	Register
1	0x0000	SCTRL: Slave mode: 64 bits words written to this address will be sent as control words. First word of an event is a control word starting with 0x5. Last word of an event is a control word starting with 0xA. The word is immediately sent to the link.
1	0x0008-0xFFFF8	SDATA: Slave mode: data addresses. The data is immediately sent to the link. Maximum size of burst is 64kB.
2	0x0000	GCTRL This register is the same as MCTRL (the same register is mapped to both addresses)
2	0x0004	GEVENT: Generate mode: 3 writes generate an event. First write is event number (bits 23..0). Second write is wordcount for payload (headers not included), in 64 bits words (bits 24..0). Third and last word is the random event generation seed (in bits 7..0) and source ID (in bits 27..16).

5.2 Software documentation (API)

The API defined for the sender board is very similar to the one of the receiver board. The software is much simpler, as it doesn't use complex structures for software/hardware synchronisation. The software currently only covers the sending of generated events.

The “open” and “close” functions are the very close to those for the receiver board. Prototypes are:

```
struct fedkit_sender * fedkit_sender_open (int index, int * status) ;
void fedkit_sender_close (struct fedkit_sender *sender) ;
```

As previously, the open function return a pointer to the receiver structure or NULL in case of failure. “status” is a pointer to the optional integer in which the error code will be stored. A value of NULL meaning “don't care” can be used for this pointer. “close” never reports failure.

5.2.1 Sending generated events

The main useful function is:

```
int fedkit_send_generated (struct fedkit_sender * sender, U16 word_count,
                           U16 seed, U32 event_trigger_number);

int fedkit_send_generated_fed_id (struct fedkit_sender * sender, U16 word_count,
                                  U16 seed, U16 fed_id, U32 event_trigger_number);
```

These functions order the sender board to generate an event with the provided parameters. They also puts the board in generator mode (as opposed to master/slave mode). The function returns a status information. At the time of writing only FK_OK (zero) or FK_overflow, if the internal FIFOs in hardware are temporarily full. FK_overflow is safe and simply means the event fragment was not sent.

The _fed_id variant, only supported by the hardware since version 0x3F200009, allows the user to set a sourceID in the headers of the event fragment. Using this function in earlier versions of the hardware simply gets the fed_id parameter to be ignored.

5.2.2 Using sender board as a slave

The PCI address for base address 1 can be retrieved with the function.

```
void * fedkit_sender_slave_bus_address (struct fedkit_sender * sender);
```

The whole range is also mapped in user space, so the user program can directly write to the board (to be used with caution).

```
void * fedkit_sender_slave_user_address (struct fedkit_sender * sender);
```

Special function should be used to provide 64 bit move/copy/write, etc... This has to be investigated (probably inline assembler). A special function that generates an event directly CPU to PCI (no memory used) has been written in the test program test_merge.cc. The main trick is to use MMX instructions movq to access 64bits of memory/PCI at once. This trick doesn't work on all machine depending on the chipset. It often generates 32 bits accesses anyhow. This leads to poor performance.

Also several strange PCI access break the functionality of the slave mode on various chipsets. Therefore the slave mode should not be used. (master mode gives better performance also).

In the future, add 2 functions: fedkit_sender_slave_bus_ctrl fedkit_sender_slave_bus_data, same for user space.

Before using the slave functions, the user has to choose the slave mode by calling:

```
void fedkit_sender_enable_slave (struct fedkit_sender * sender);
```

5.2.3 Using sender board as master

To use the board as a master, the user just has to call:

```
int fedkit_master_send (struct fedkit_sender * sender, U32 pciaddress,
                        U16 word count);
```

The DMA buffers (Dbuff) provided by the xdaq-shell can be used in this case. The master/slave mode (in the hardware) is automatically selected by this function.

`fedkit_master_send` returns FK_OK on success. Any other return value means that the hardware or software queues was NOT able to store the block address. The user can safely retry to send the same block on this error.

As with version 1_18 of the software, a new scheme using the DMA done interrupt for master mode has been introduced. In this case, several blocks of data are stored in a fixed size ring buffer. The

fedkit_master_send function only pushed a given block in the ring buffer. Therefore, before re-using a block that was previously used in this FIFO, the user should make sure that this block was transmitted.

The following function returns non zero if the given address is not any more in the queue, and is not being processed by the DMA engine (and therefore the block can be reused) and zero if the blocks is still in DMA or in the queue.

```
int fedkit_master_send_complete (struct fedkit_Sender * sender, U32 pciaddress);
```

5.2.4 Getting the FPGA version

As for the receiver, the sender's firm ware revision number can be read with:

```
U32 fedkit_sender_get_FPGA_version (struct fedkit_sender * sender);
```

6 Test programs

6.1 ‘test_merge’: FEDkit’s main test program

During the development of the FEDkit, a test program was written to test all the features of the FEDkit. It was also used to measure the performance and the reliability. This program was developed for internal purposes only (debugging during development), but it can be a useful tool for quick checks. The default behaviour test_merge is to handle sending and receiving in the same computer. This can be realized either by using a self sender board (for software debugging) or by using one or two senders and a receiver in the same PC. Split operation is also possible, by running two instances of test_merge, one sending, one receiving.

All the tunable parameters of the fedkit are accessible from the command line. The sender boards can operate in generation, master or slave mode.

Finally, the test_merge program can optional check all the data received to validate the received data.

When the receiver is a self sender hardware, the program detects this automatically, and no sender board is looked for.

The online help of the program reads as follows:

```
./test_merge --help
./test_merge: tests a fedkit 1-or-2-senders-receiver setup (all in the same machine)
Usage : ./test_merge [-s <sender unit number>] [-t <second sender unit number>]
[-r <receiver unit number>] [-f <fixed word count>] [-M <max word count>] [-m <min word count>] [-I 1]
[-n <fragment number>] [-c {0|1|2}] [-L <sender pattern>] [-o {0|1}] [-a {0|1}] [-T {0|1}]
[-b <block number>] [-B <block size>] [-h <header size>] [-s {0|1}] [-R {0|1}] [-K {0|1}] [-d <send delay in usecs>]
[-i <initial seed>] [-F <trigger/fed number>] [-J <just receive/send mode>]
[--fixed-seed <fixed seed value>] [--daq-column] [--no-MMX] [--maxperf]

-s, -t, -r indexes of sender unit(s) (1st and second) and to the receiver unit
    defaults are -s 0 -t 1 -r 0
-f, or -M and -m defines the sizes of the fragments (fixed size, max and min)
    defaults are -m 10 -M 10000 (using a fixed size in the command line excludes random sizes)
-I if set to non-zero, this parameters instructs the program to use an incremental word count
-c turns on and off the event payload checking
    default is off (-c 0), 1 is check all by CRC/truncated flags, 2 check only structure, 3 is all including flags
-n number of event fragments; default is -n 50
```

```

-L sender pattern. Pattern tells the receiver from which input it should expect event
fragments. 1 means only from first input, 2 only from second input, 3 from both
default is -L 3
-o turns on an off the final test on open close. Tries 100's of open-closes in sequence; default is off
-a turns on and off the allocation of blocks by the fedkit (0=noalloc, 1=normal); defaults is on
-T turns on and off the multithreaded test; default is off
-b block number; default is 2000
-B block size. Has to be 64bits aligned. Will be ignored if not. Default is 4k
-h header size. Has to be 64bits aligned. Will be ignored if not. Defaults to 0
-S slow sender. Setting this flag to 1 will enable emulation of a slow sender. Default is off.
-R slow receiver. Setting this flag to 1 will enable emulation of a slow receiver. Default is off.
-d sender delay. This value (by default 0) is the time the sender thread should wait after sending a fragment.
-i initial seed. Set the seed value at the beginning of the program. Default is 0.
-F initial trigger number. Set the initial value for the (fake) trigger number. Default is 1000.
-g event generation mode. Choose between master, slave, and hardware generated events.
    default is 0 (hardware), 1 is slave mode, 2 is master mode.
-K turns on and off benchmarking. (Checks several performances on the PC).
-J chooses to Just send/or receive events. 1 is just send. 2 is just receive. If check is enabled
    the receive process will guess the parameters the same way as the send process. The delays
    for the receiver are increased to infinite for the first event fragment. ctrl-c aborts that.
--daq-column set up everything so that we can run with the daq column prototype.
--fixed-seed indicates a fixed value for the seed. Seed normally increments during run.
--no-MMX hints the program not to use MMX. Currently only not using MMX to access PCI.
--maxperf Tell test_merge to maximize the performance, i.e. to minimize the checks

```

-s, -t and -r indicate which sender(s) and receivers to use in systems where there are several boards installed. The fedkit can handle 8 senders and 8 receivers in the same system.

Besides setting fragments parameters like size or seed, the most important options are -L, -J, and -c.

-L indicates which channels will be used in the test. The fedkit can use channel 0, 1 or both on the receiver. The fedkit assumes the first sender is connected to channel 0 and the second is connected to channel 1. The value of this parameter is a bit mask for the channel, so '-L 1' means channel 0, '-L 2' means channel 1 and '-L 3' means both channels. When using only channel 1, the second sender is the one that will be used (so option '-s' is useless in this case, and it's likely than one wants to use '-t 0' if there is only one sender in the system).

-J allows split operation of test merge, with one instance sending (-J 1) and one instance receiving (-J 2). If the checking option is used, the receiver side has to know the generation parameters so options like -f, -m, -M, --fixed-seed, -I, -n, -i, -F have to be passed the same way to both sender and receiver. If the receiver is not checking on data, the sizes of received fragments are added and the pay loads and headers are not looked at. In any case, calling both instances with the same parameters in both sides (with just -J option differing) is safe.

'-c 1' enables checking of data. This option relies on the payload being generated with a pseudo-random generator. The software generated modes (master mode and slave mode) uses the same generation algorithm as the hardware generation mode does, so that there is only a single checking loop. Checking limits the bandwidth to roughly 200MB/s. *Is it? Need measurements!*

6.2 ‘fedkit_dump_receiver’: the link level debugging tool

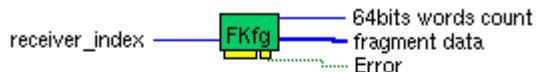
The fedkit package also comes with a program that dumps to the terminal the words coming from one S-LINK64 interface. The default behaviour is to dump link 0 of receiver 0, but the command line options allow to choose the receiver (with '-r') and the line number (0 or 1) with '-l'.



7 Labview VIs

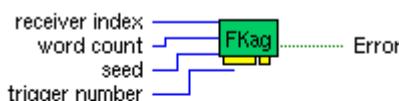
Some functions of the fedkit API are now wrapped in LAbView VIs. The VIs were designed on Linux version of LAbView 6.1.

Those VIs are:



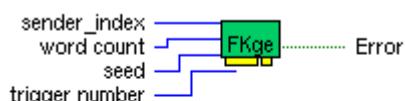
fedkit_frag_get.vi

This VI retries a fragment from a FEDkit receiver and stores the data in an unsigned 32 bits integer array. Output also includes word count and an error flag.



fedkit_auto_generate.vi

This VI orders the self sender receiver to generate an event fragment (pseudo-random generation).



fedkit_generate.vi

This VI orders the sender to generate an event fragment (pseudo-random generator).

The VIs can be found in the “.../itools/packages/fedkit” directory after doing a ‘make fedkit-lv’.

8 Going further

The FEDkit is just a part of the software packages provided by the DAQ group. Other packages can be used to merge data from several sources, set parameters, use data bases and managers as in the DAQ run control. The FED board is in this case automatically configured into the executive software. This configuration can be stored in a database. Several FED boards can be used, their data merged, monitored. The whole package and documentation can be found on the XDAQ web page [5].

References

- 1 Readout unit working group web page on CMS DAQ horizontal web page
<http://cmsdoc.cern.ch/cms/TRIDAS/horizontal/>
- 2 S-link 64 specification <http://cmsdoc.cern.ch/cms/TRIDAS/horizontal/docs/slink64.pdf>
- 3 FRL documentation <http://cmsdoc.cern.ch/cms/TRIDAS/html/Documents.html>
- 4 i2ocore documentation <http://cern.ch/cano/i2ocore>
- 5 XDAQ web page <http://cern.ch/xdaq>
- 6 Fedkit web page <http://cern.ch/cano/fedkit>

